Network Addresses in NNDK 2.8+

Dan Ciliske

San Diego, CA
Netburner, Inc.

May 9, 2016

Overview

Sections:

- Adressing Modes
- How it works
- How it goes wrong

Addressing Modes

- IPv4
  - Static
  - DHCP
  - AutoIP
- IPv6
  - Static
  - AutoConfig
  - DHCPv6

IPv4: Static

Pros:

- Easy to track what address is assigned to device.
- Guaranteed to have address.

Cons:

- **Must** assign address to device when provisioning.
- Chance of collision if used on network with DHCP.

NetBurner
Networking in One Day!

IPv4: DHCP

Pros:

- Automatically assigns IP from central server.
- Also provides DNS, Gateway, and other information.
- Possible to fluidly and automatically update DNS info.

Cons:

- **Requires** central DHCP server. If none exists, fails to obtain address.
- Requires device negotiate address. This takes time and generates broadcast messages.

NetBurner
Networking in One Day!

IPv4: AutoIP

Pros:

- Device automagically assigns address to itself.
- Builtin process for discovery of address conflicts and automatic reassignment.
- Can have a predictable address sequence to try.

Cons:

- Can only be used for link local communications.
- Does not assign any additional information (DNS Servers, Gateway, etc.)
- PCs can take a long time upon first connection to allow access to AutoIP subnet.

NetBurner
Networking in One Day!

IPv6: Static

Pros:

- Easy to track what address is assigned to device.
- Guaranteed to have address.

Cons:

- **Must** assign address to device when provisioning.
- Chance of collision if used on network with DHCP or AutoConfig. (AutoConfig has no reserved subnet in IPv6.)

NetBurner
Networking in One Day!

IPv6: AutoConfig

Pros:

- Device automagically assigns address to itself.
- Builtin process for discovery of address conflicts.
- All addresses share same suffix related to device MAC.
- Link-local address required for all other IPv6 use.

Cons:

- Requires router to set AutoConfig bit to config globally routeable address.
- Router may not provide DNS server address.

NetBurner
Networking in One Day!

IPv6: DHCPv6

Pros:

- Automatically assigns IP from central server.
- Also provides DNS, Gateway, and other information.
- Possible to fluidly and automatically update DNS info.

Cons:

- **Requires** central DHCP server. If none exists, fails to obtain address.
- Requires device negotiate address. This takes time.

NetBurner
Networking in One Day!

IPv6: Other Notes

Other Notes:

- A managed DHCPv6 server should never be used for an address space (i.e. prefix) that any router on the link provides with the Autoconfig bit set.
- DHCPv6 can also be used simply for "Additional Information" without assigning an address.
- IPv6 is definitionally a multihomed system. A device will have multiple addresses per interface in most real world scenarios.

NetBurner
Networking in One Day!

How it works

How do we go from boot to communicating over IPv6?

NetBurner
Networking in One Day!

Hello World

- Boot the module.
- InitializeStack called to start Network Stack.
- Module starts establishing AutoConfigured Link Local IPv6 address. (fe80::. . . , uses reserved prefix akin to 192.168/16)
- Module determines address as a function of device MAC.
  - Link local is definitionally autoconfigured.
- Performs Duplicate Address Detection (DAD) to confirm address not in use.
  - If it's clear, we use it.
  - If it's in use, pack our bags, we're done. We must have LinkLocal address to do any IPv6. This is a hard failure condition that should never happen.

NetBurner
Networking in One Day!

Getting LinkLocal

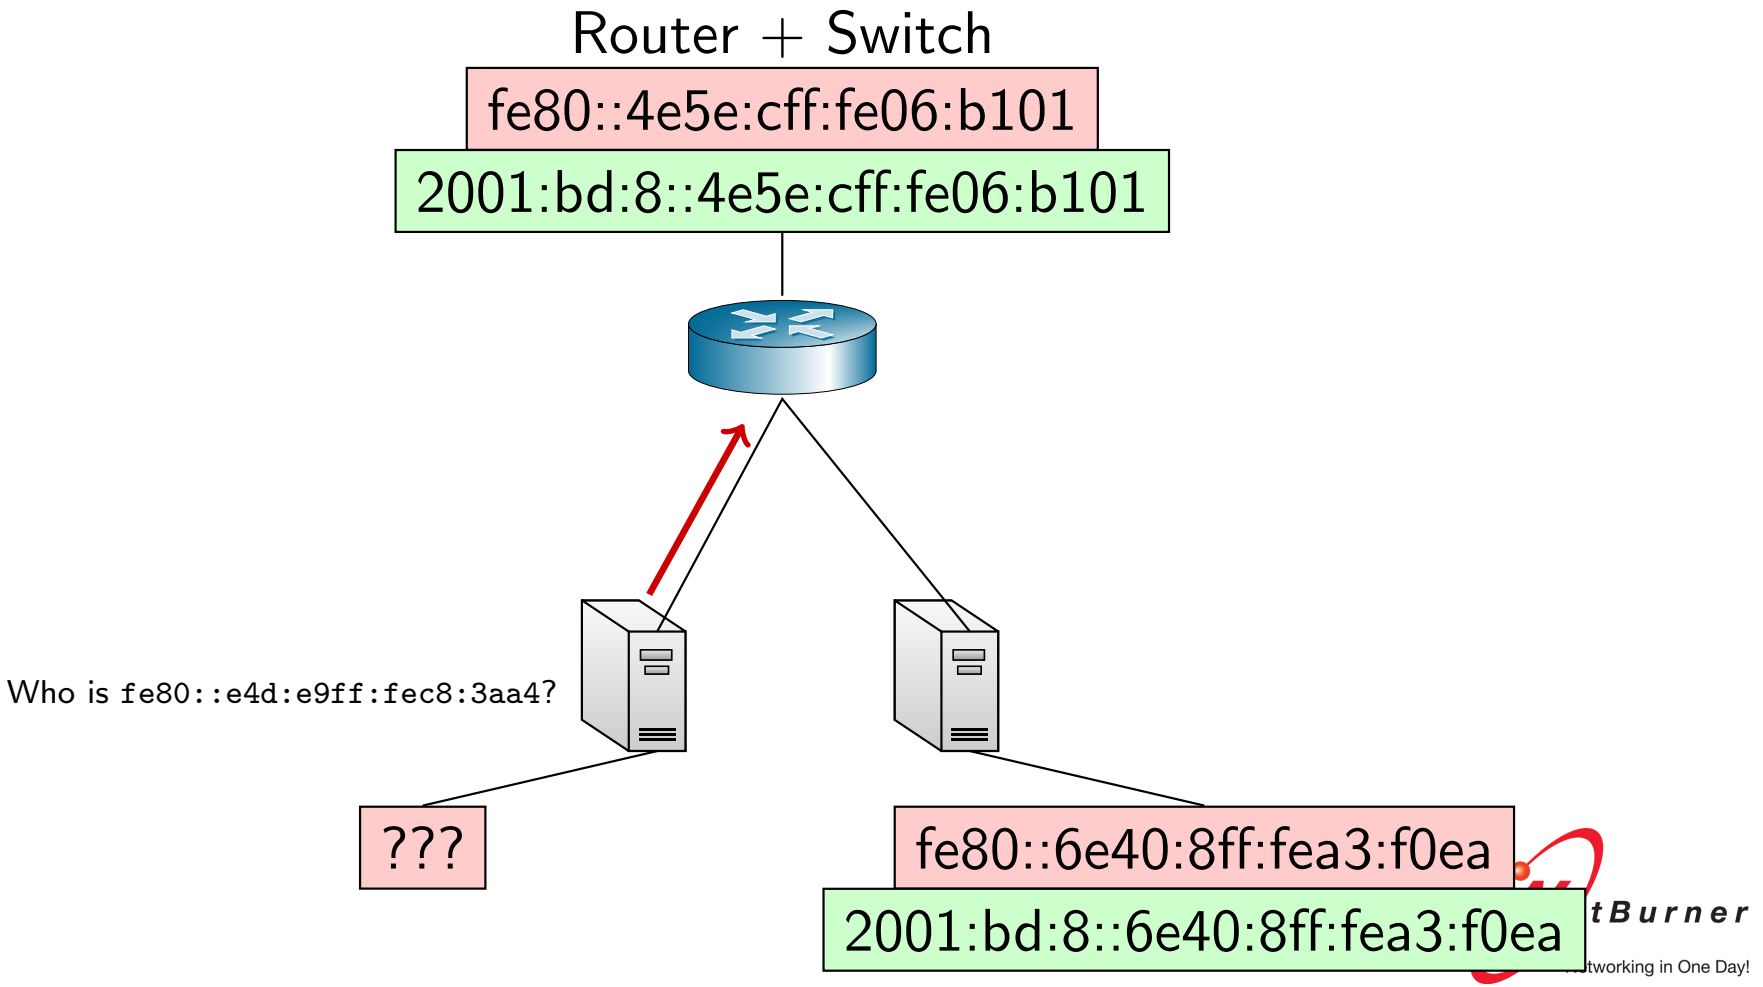

Getting LinkLocal

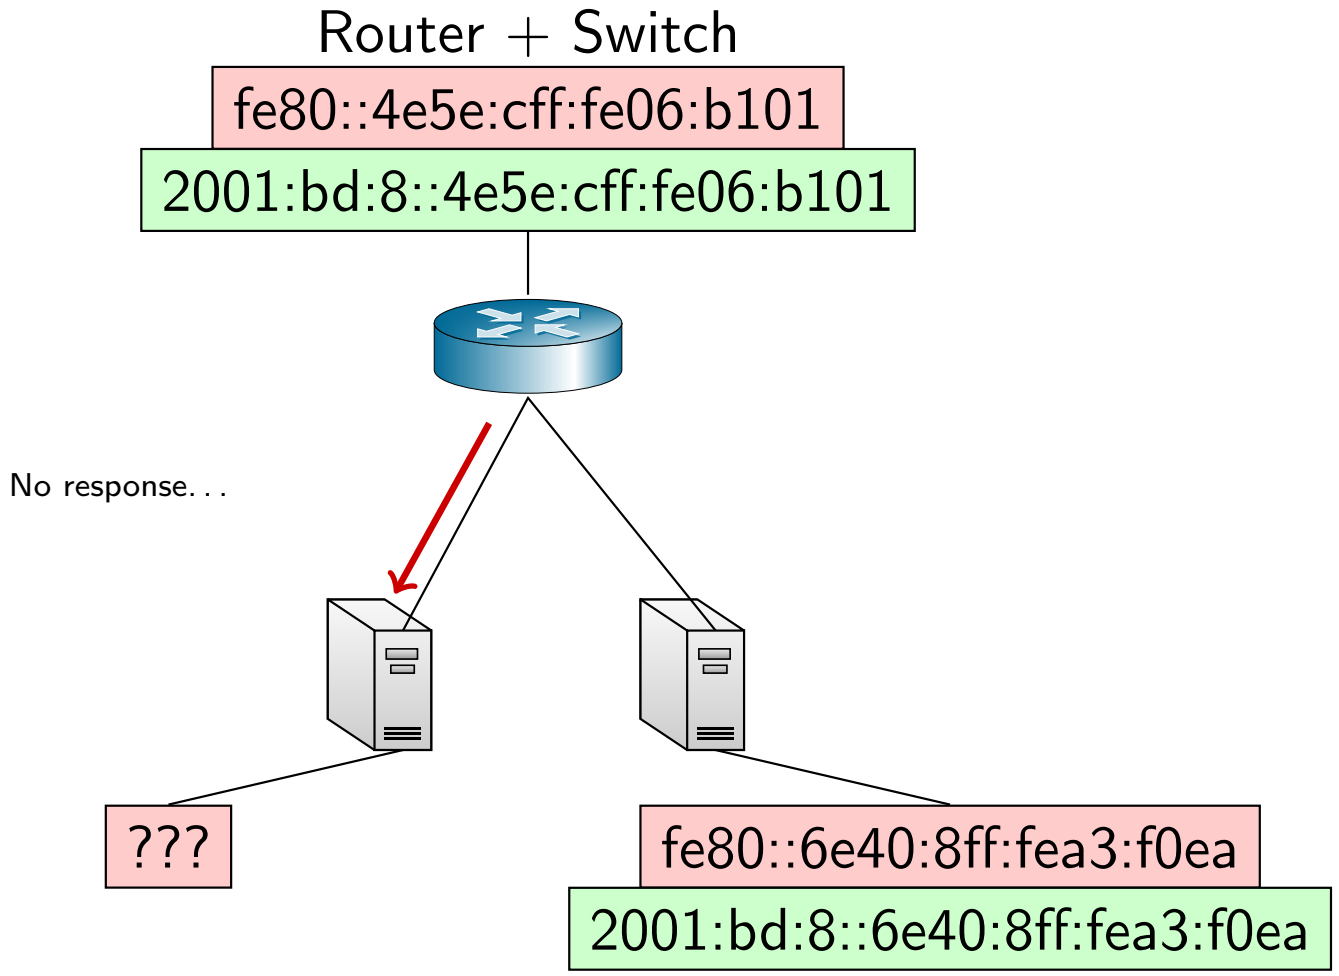

NetBurner
Networking in One Day!

Getting LinkLocal

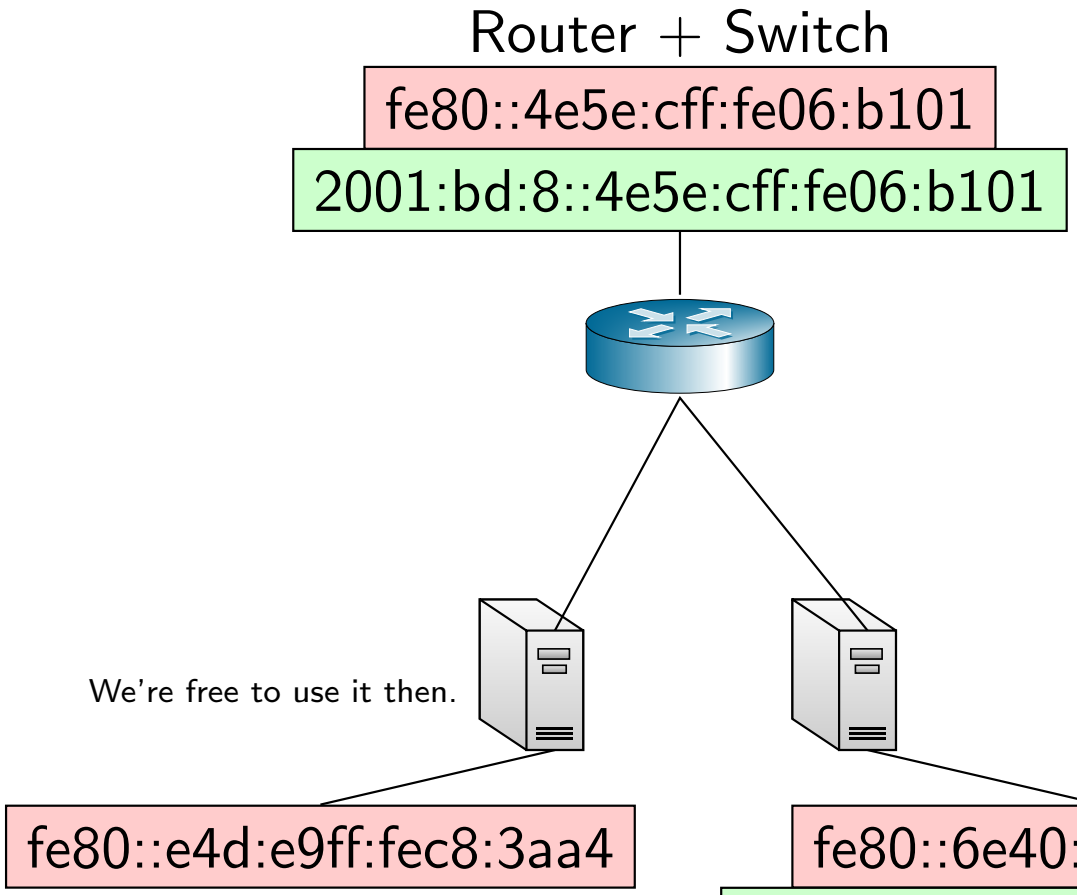

NetBurner
Networking in One Day!

I'm FE80, who are you?

- Now that the module has a Link Local address we start exploring.
- The goal is to end up with at least one globally routeable address.
- Step one towards that is to look for a router on the link.
- Module performs router solicitation. If there's a router, it sends back an advertisement. All routers listen to `ff02::2`.
- If the router can route, the advert contains a prefix designator.

NetBurner
Networking in One Day!

Processing the advert

- The advert can also contain three specific bits: the managed bit, other bit, and the autoconfig bit of the prefix option.
- These bits describe the negotiation processes to be employed.
  - Managed: There is a DHCPv6 server on the link that will perform address assignment.
  - Other: There is a DHCPv6 server on the link that will provide "additional" information.
  - Autoconfig: The advertised prefix is valid to use as part of an AutoConfigured address.

Begin Autoconfig

- If the Autoconfig bit is set, we pick a new address in the same way we chose our Link Local address, but with a different prefix
  - i.e. Say we receive prefix '2001:db:8::'. Our new address would be the Link Local address, but with the 'fe80::' replaced with '2001:db:8::'.
- After selecting the new address, we perform Duplicate Address Detection to confirm it's available.
  - If it's in use, we're done with it. We do not assign it for use.
  - If it's not in use, we're free to use it and we mark it as such.

NetBurner
Networking in One Day!

Getting Global: Autoconfig

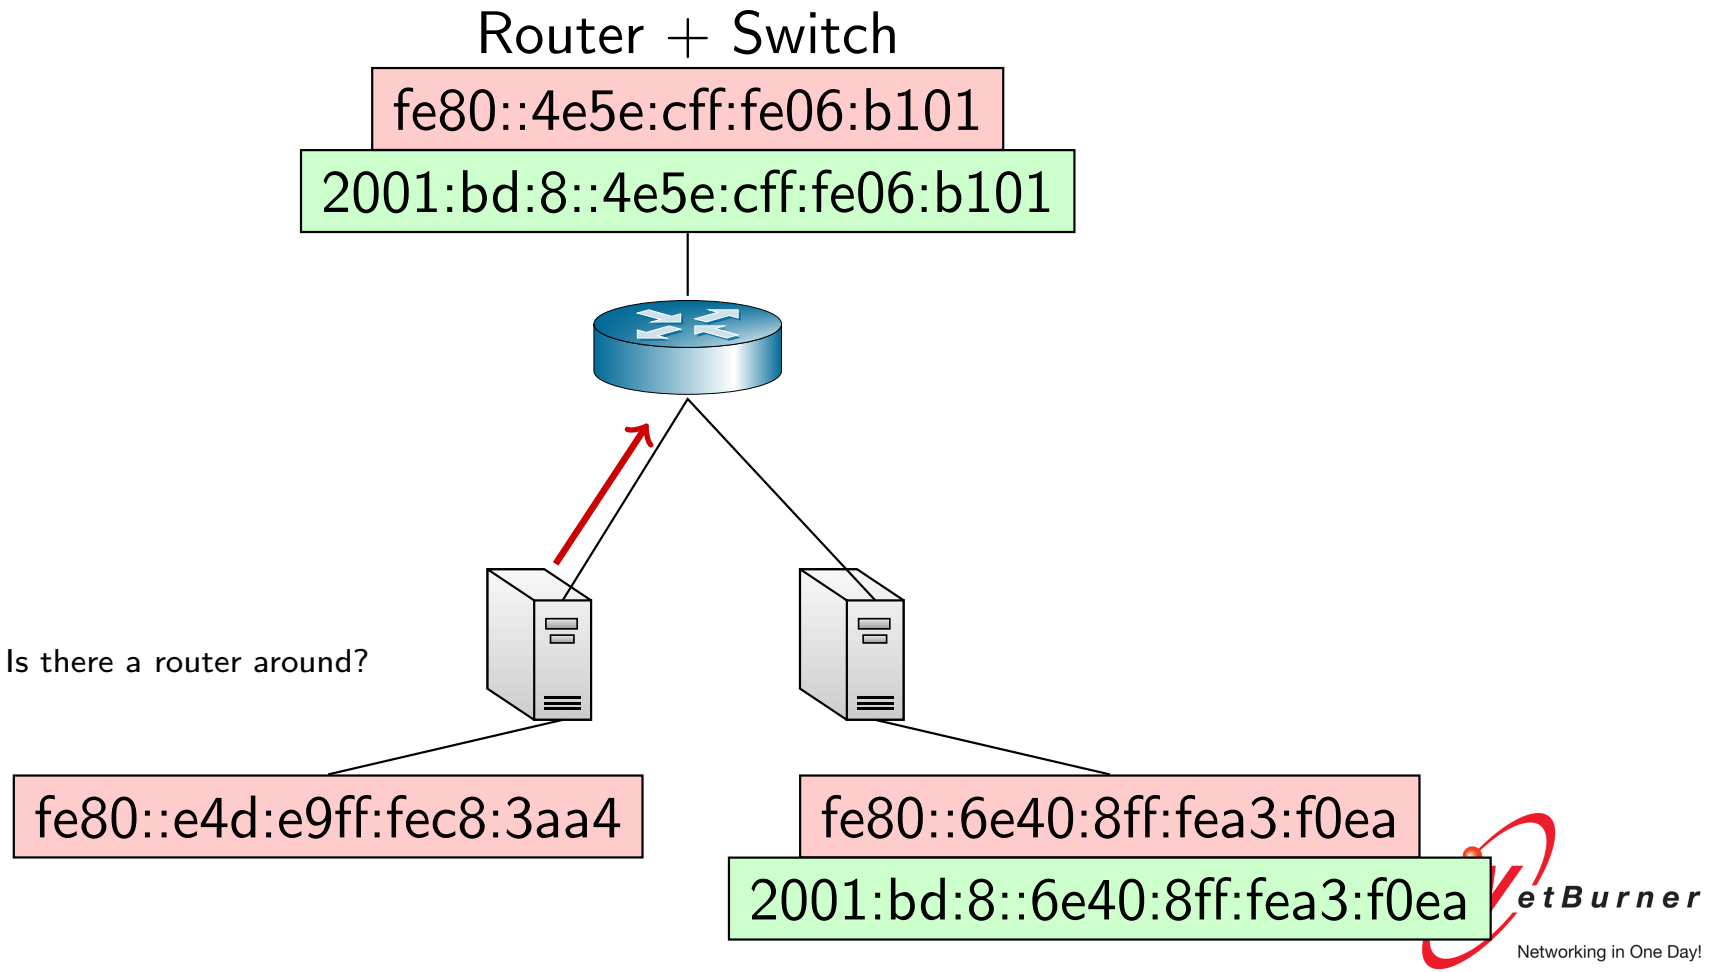

Getting Global: Autoconfig

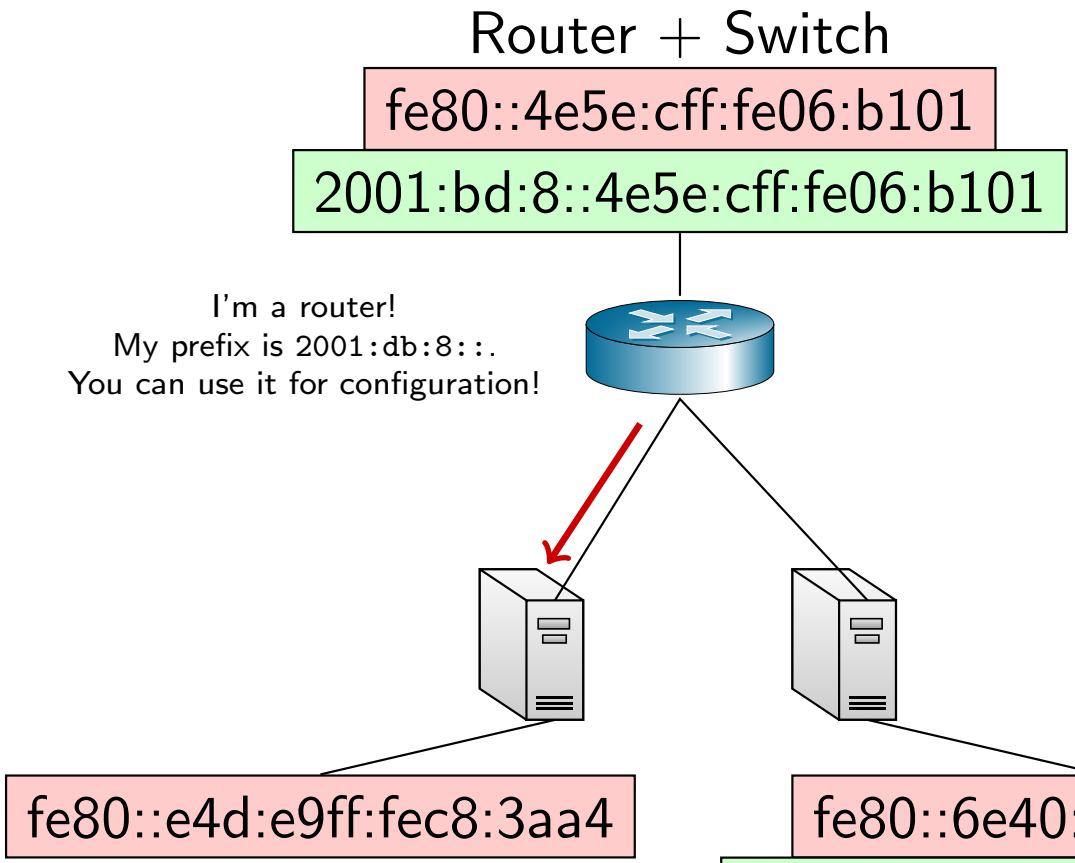

Getting Global: Autoconfig

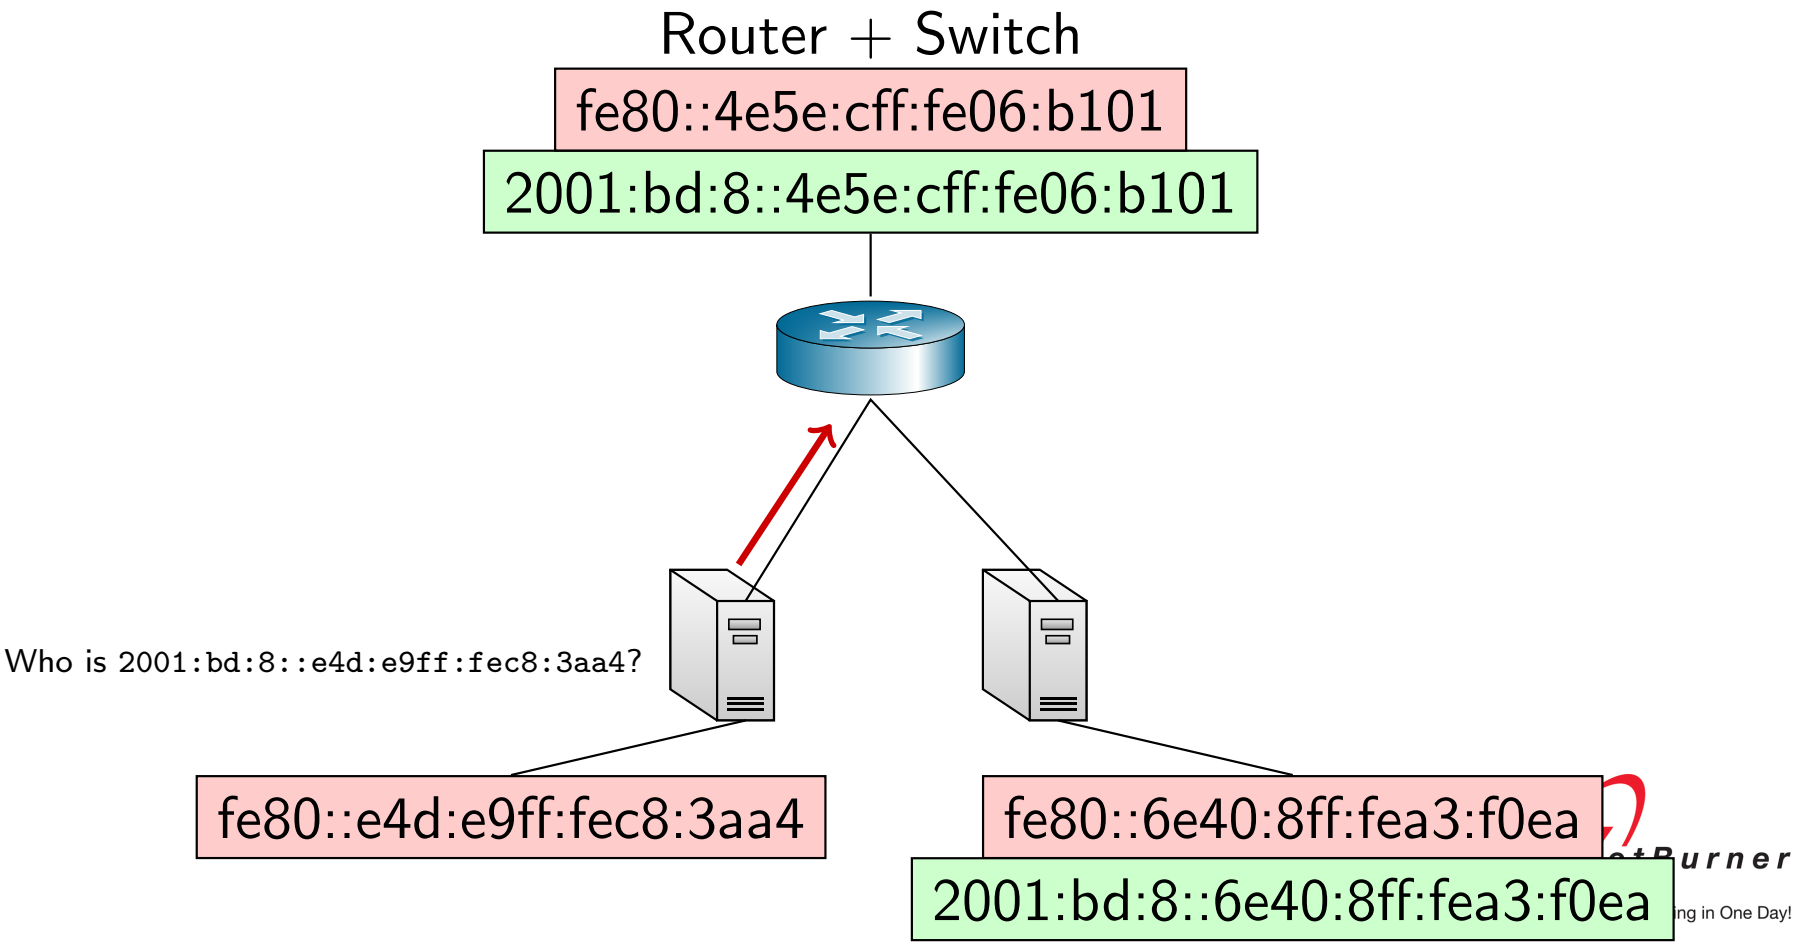

Getting Global: Autoconfig

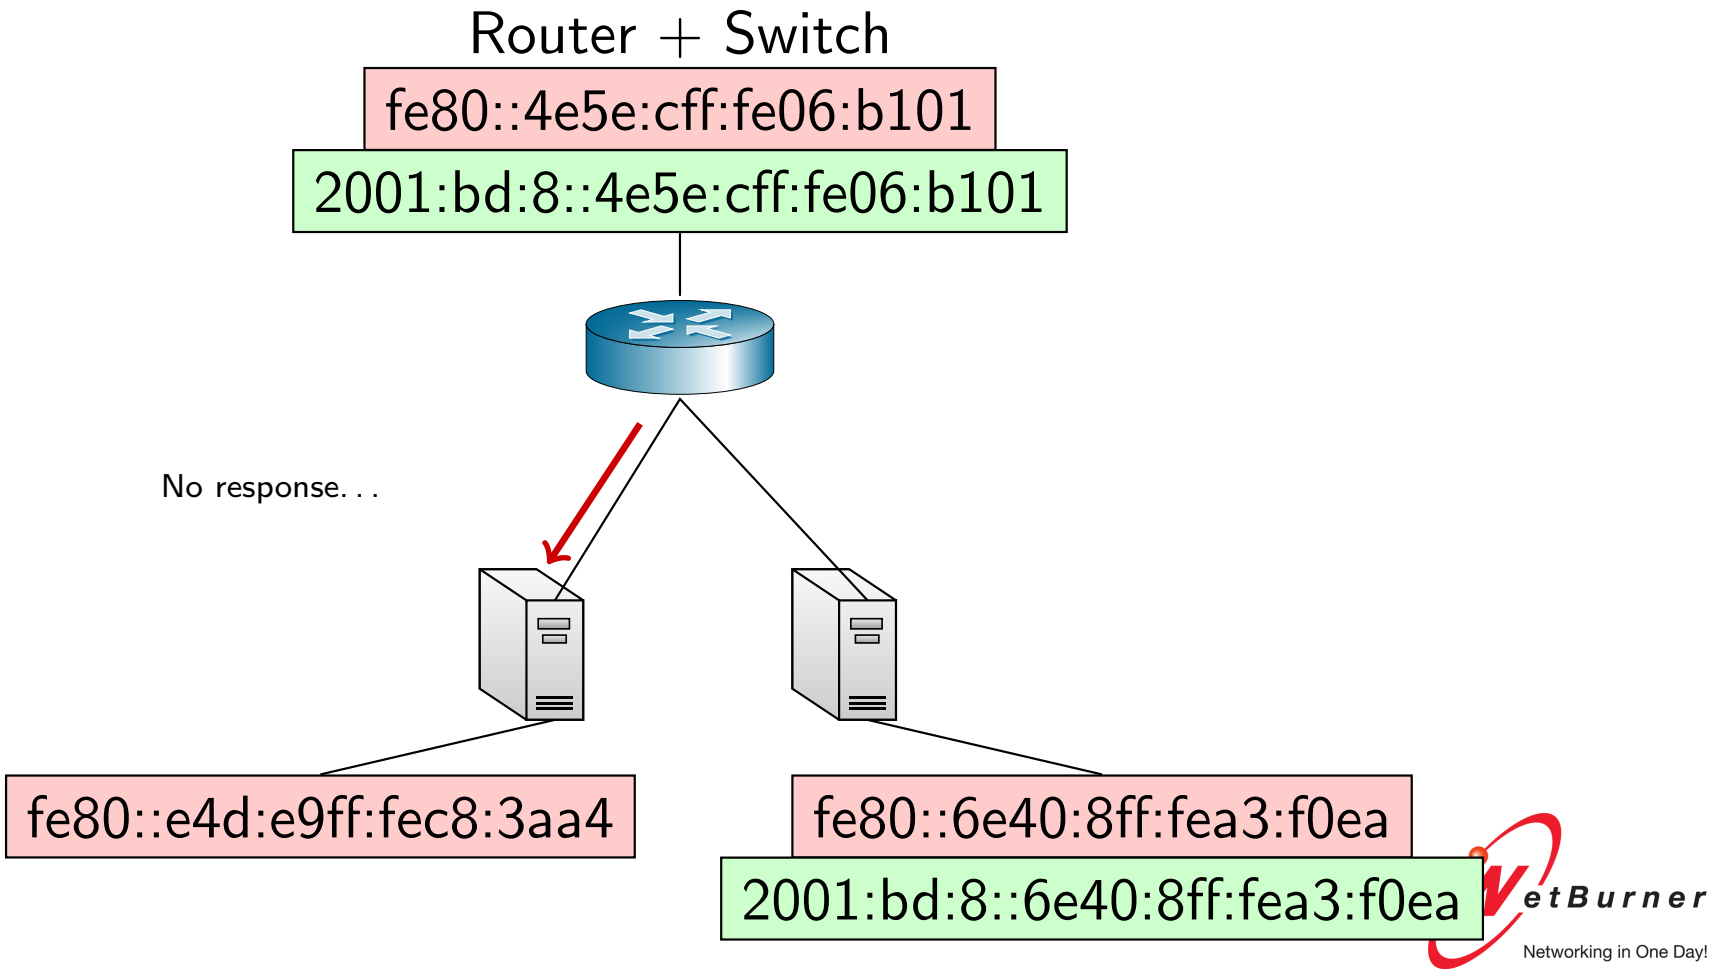

Getting Global: Autoconfig

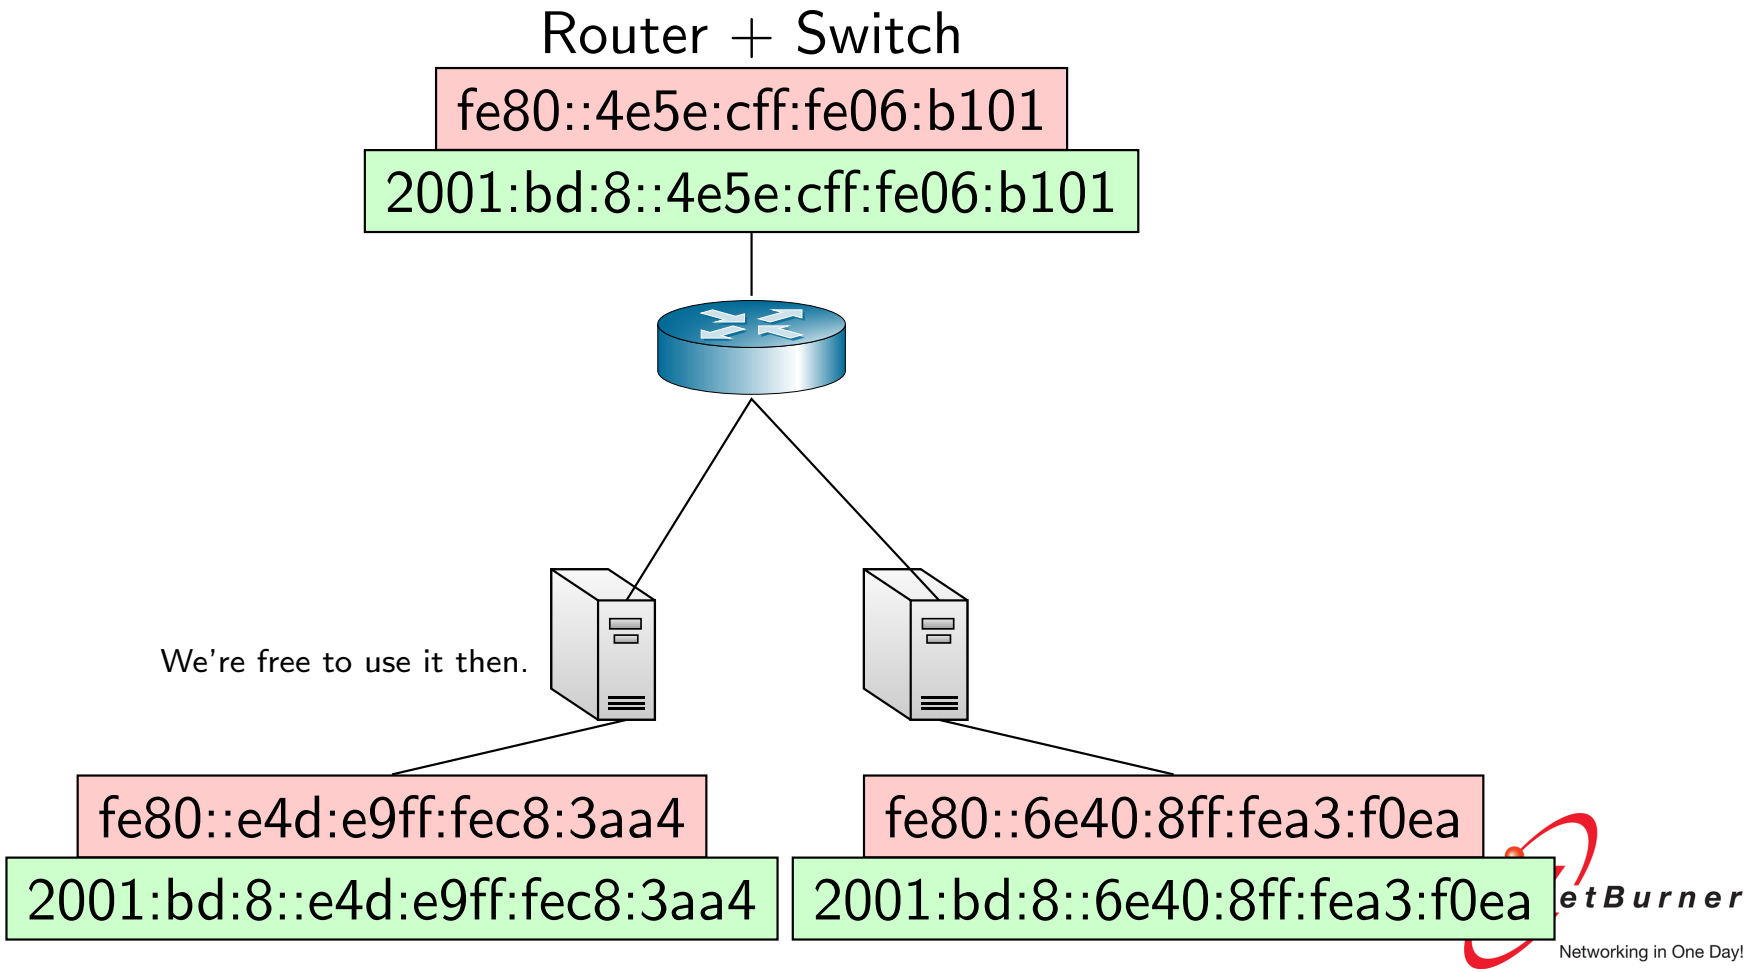

Enter DHCPv6

There's two types of DHCPv6 behaviours, based on which bits the router sets:

- If the managed bit is set, there exists a DHCPv6 server on the link that will assign addresses.
  - In this case we will perform full address negotiation.
  - This involves: client solicit, server advert, client request, then server response.
  - The DHCPv6 negotiation can also provide additional information such as DNS or NTP Servers. We request and process any DNS server responses.
  - All DHCPv6 servers and relays listen on `ff02::1:2`.

Enter DHCPv6

There's two types of DHCPv6 behaviours, based on which bits the router sets:

- If the managed bit is not set, but the Other bit is: there exists an informational DHCPv6 server on the link.
  - In this case we will only perform an Information Request.
  - This involves: client request, then server response.
  - We only request DNS server info at this point in time.
  - The client specifiies exactly what information it wants.

Getting Global: DHCPv6

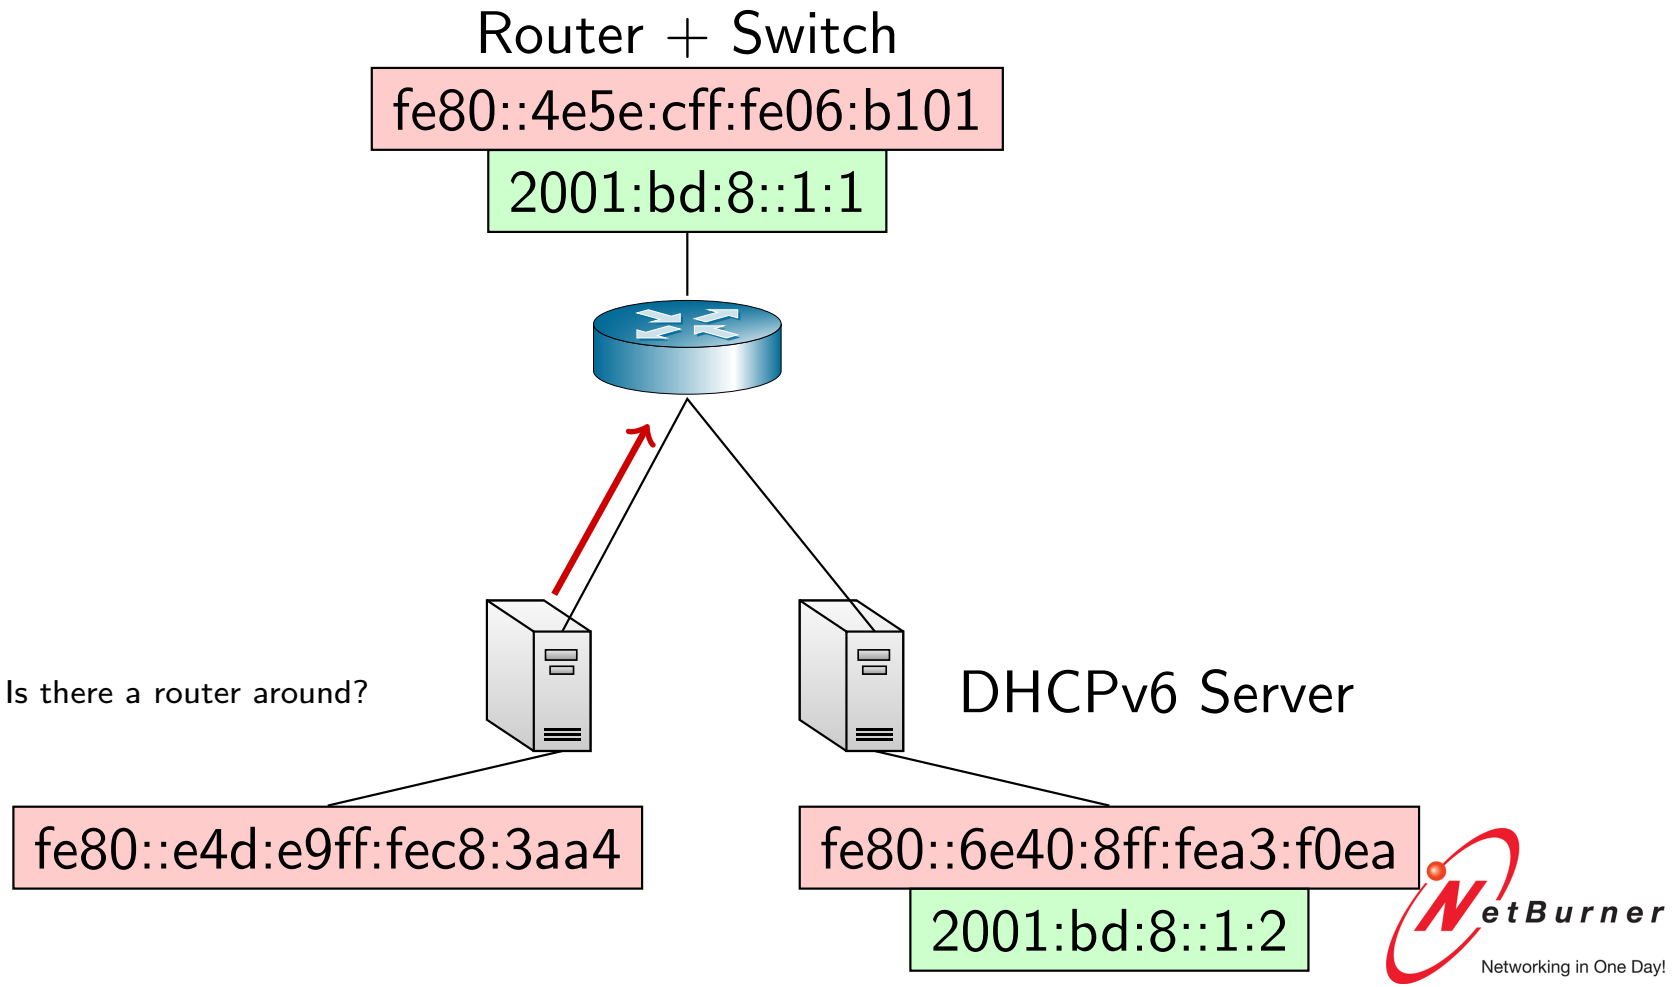

NetBurner
Networking in One Day!

Getting Global: DHCPv6

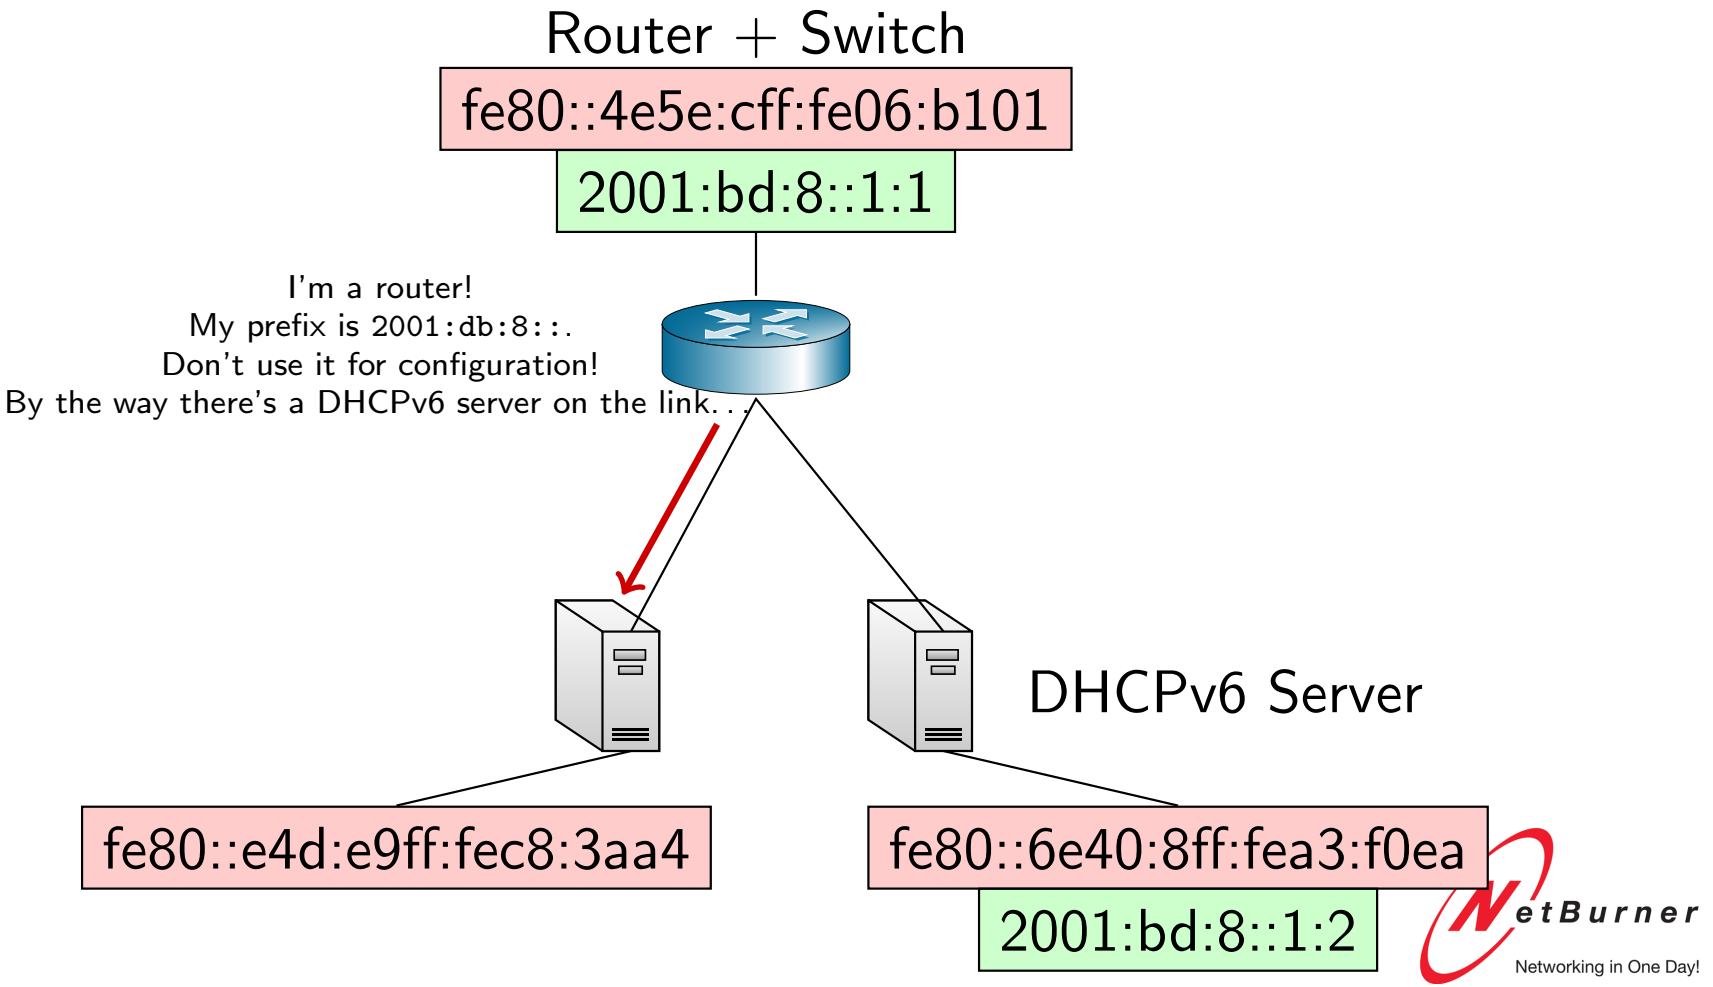

NetBurner
Networking in One Day!

Getting Global: DHCPv6

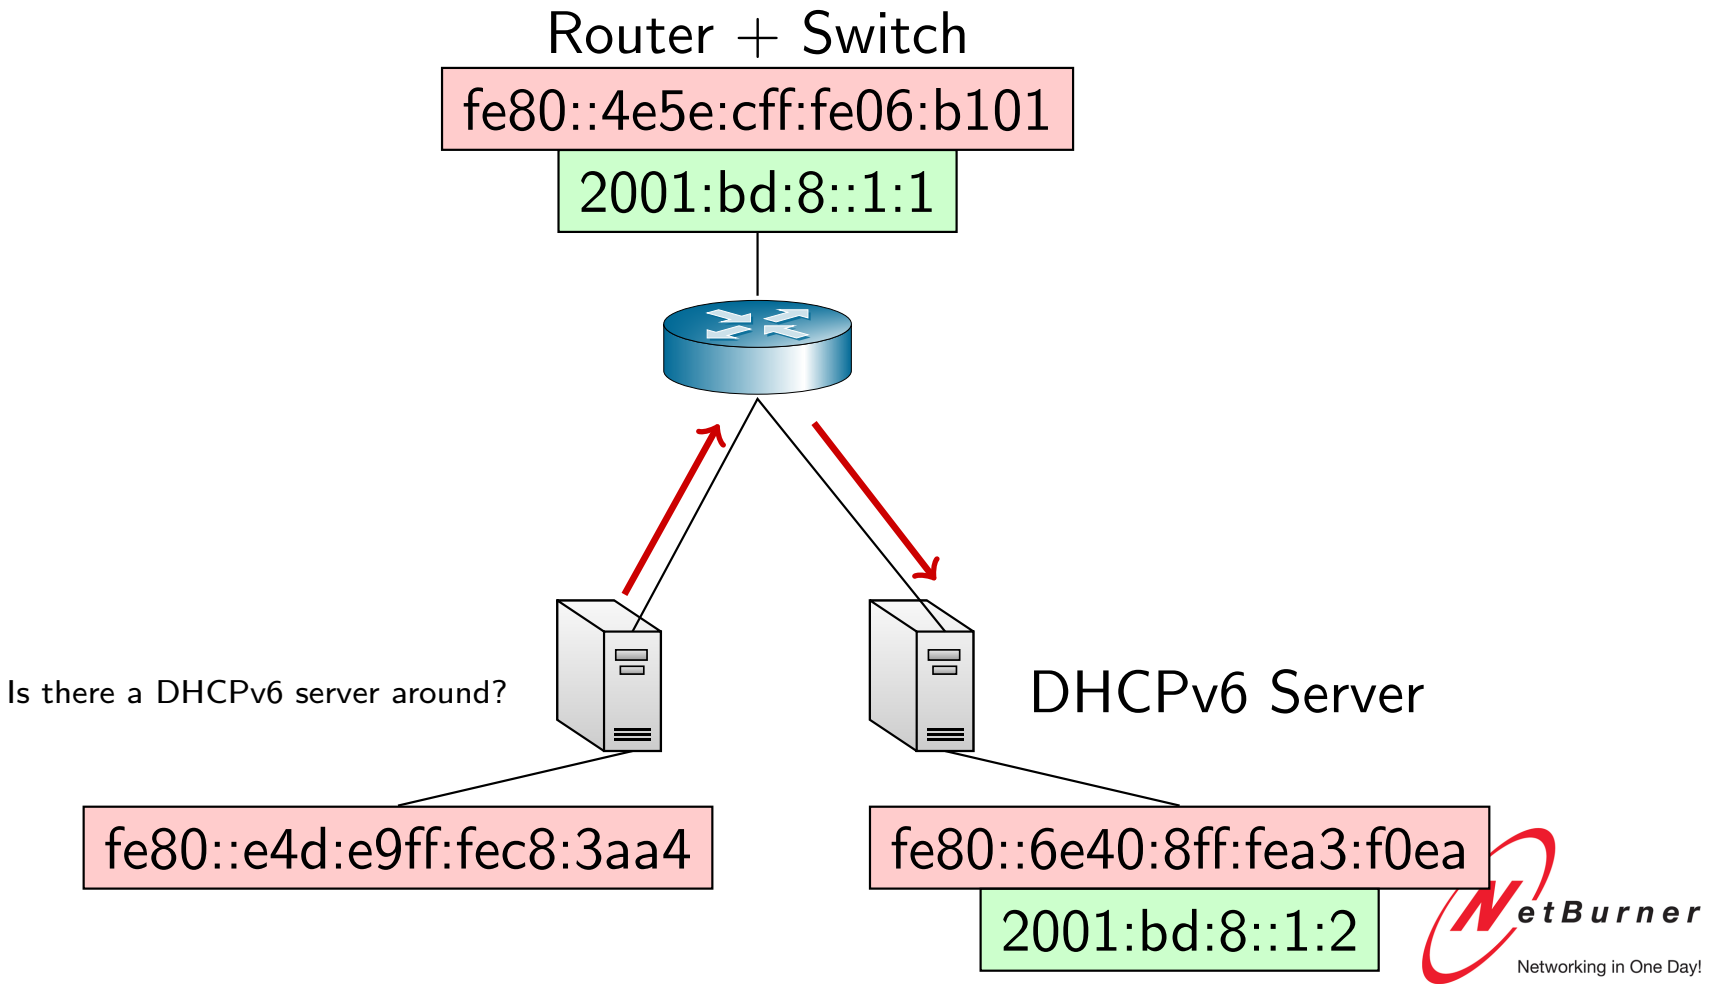

Getting Global: DHCPv6

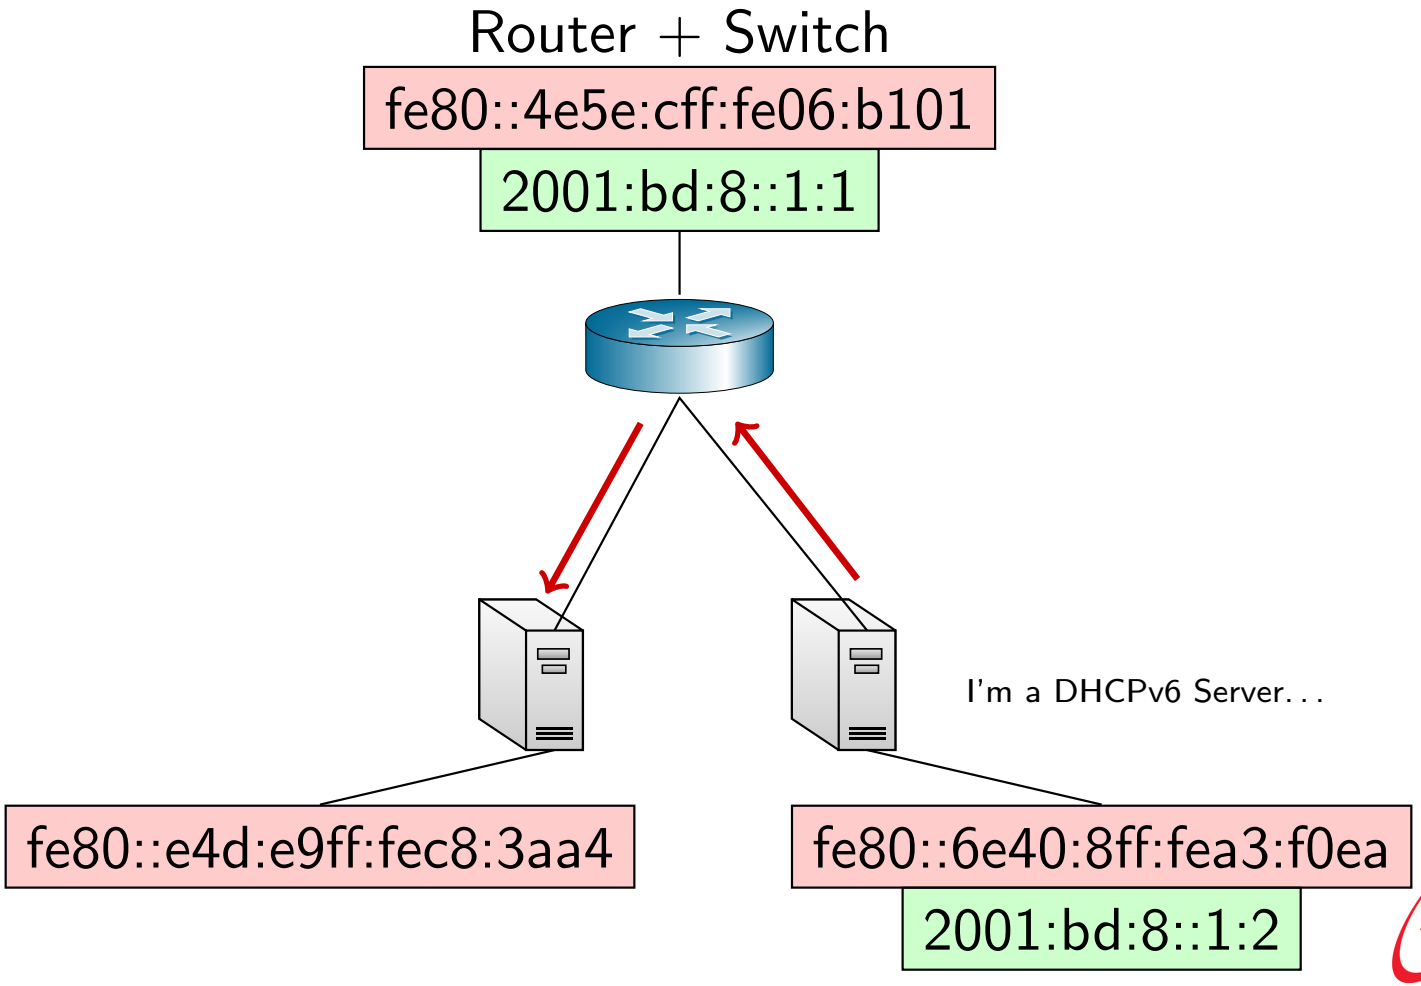

NetBurner
Networking in One Day!

Getting Global: DHCPv6

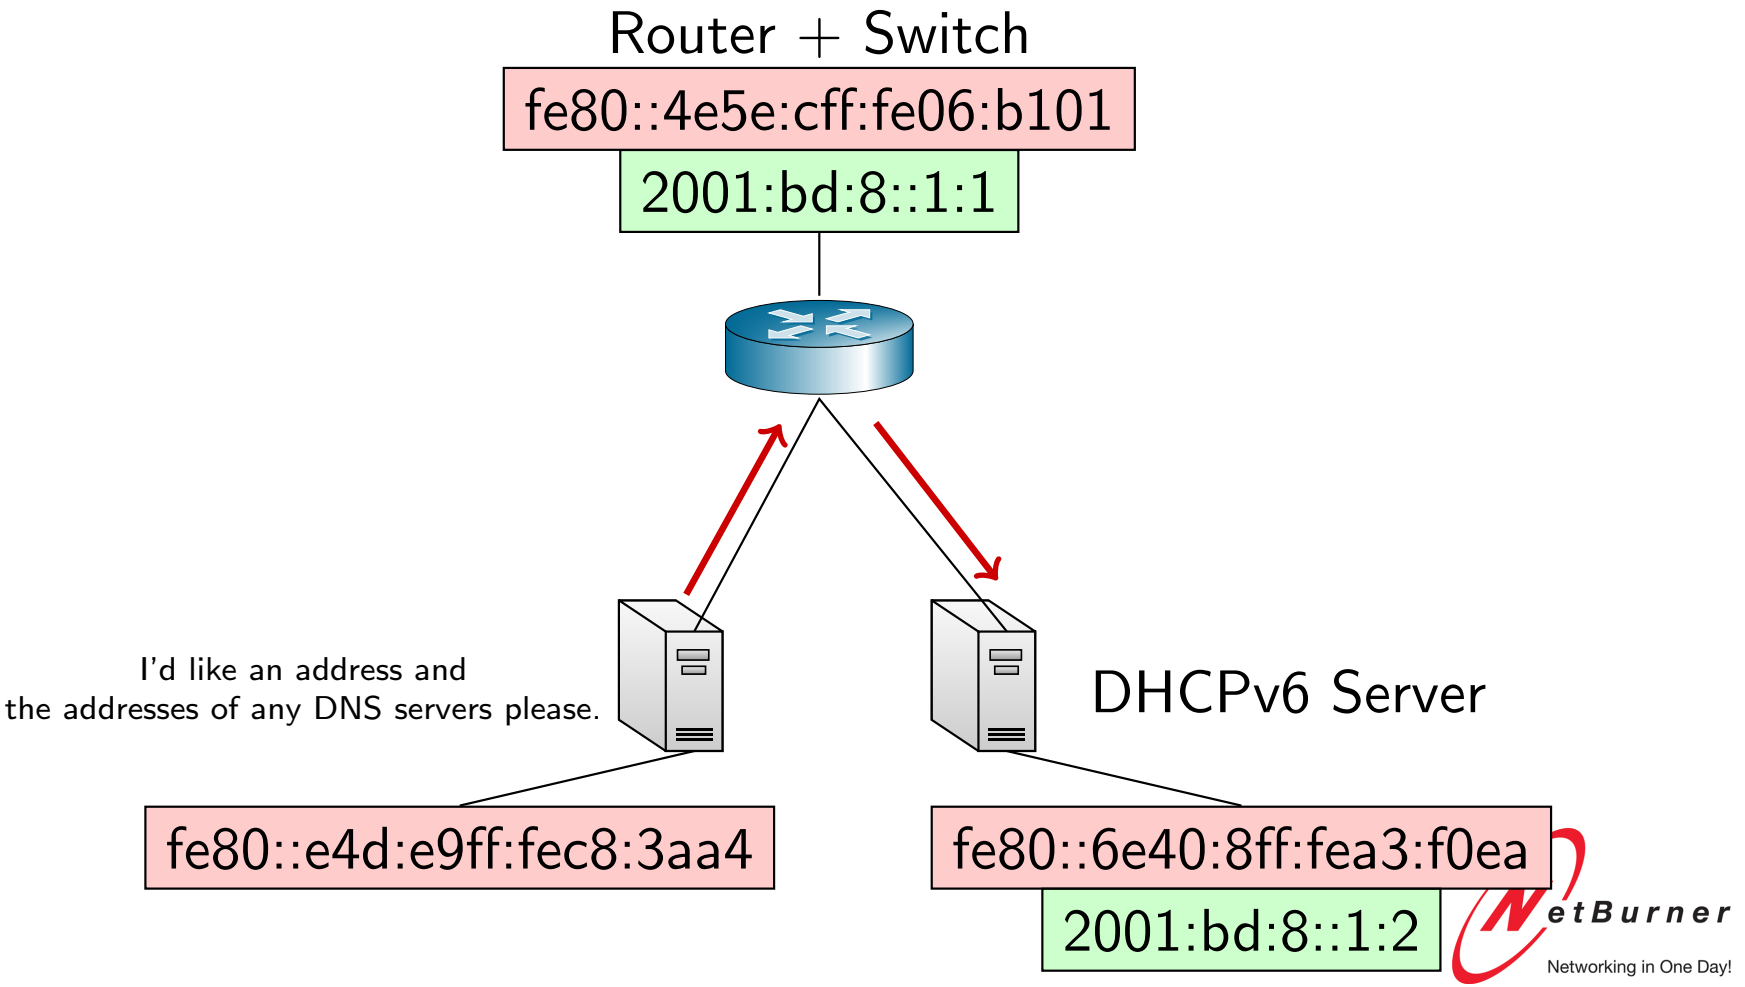

Getting Global: DHCPv6

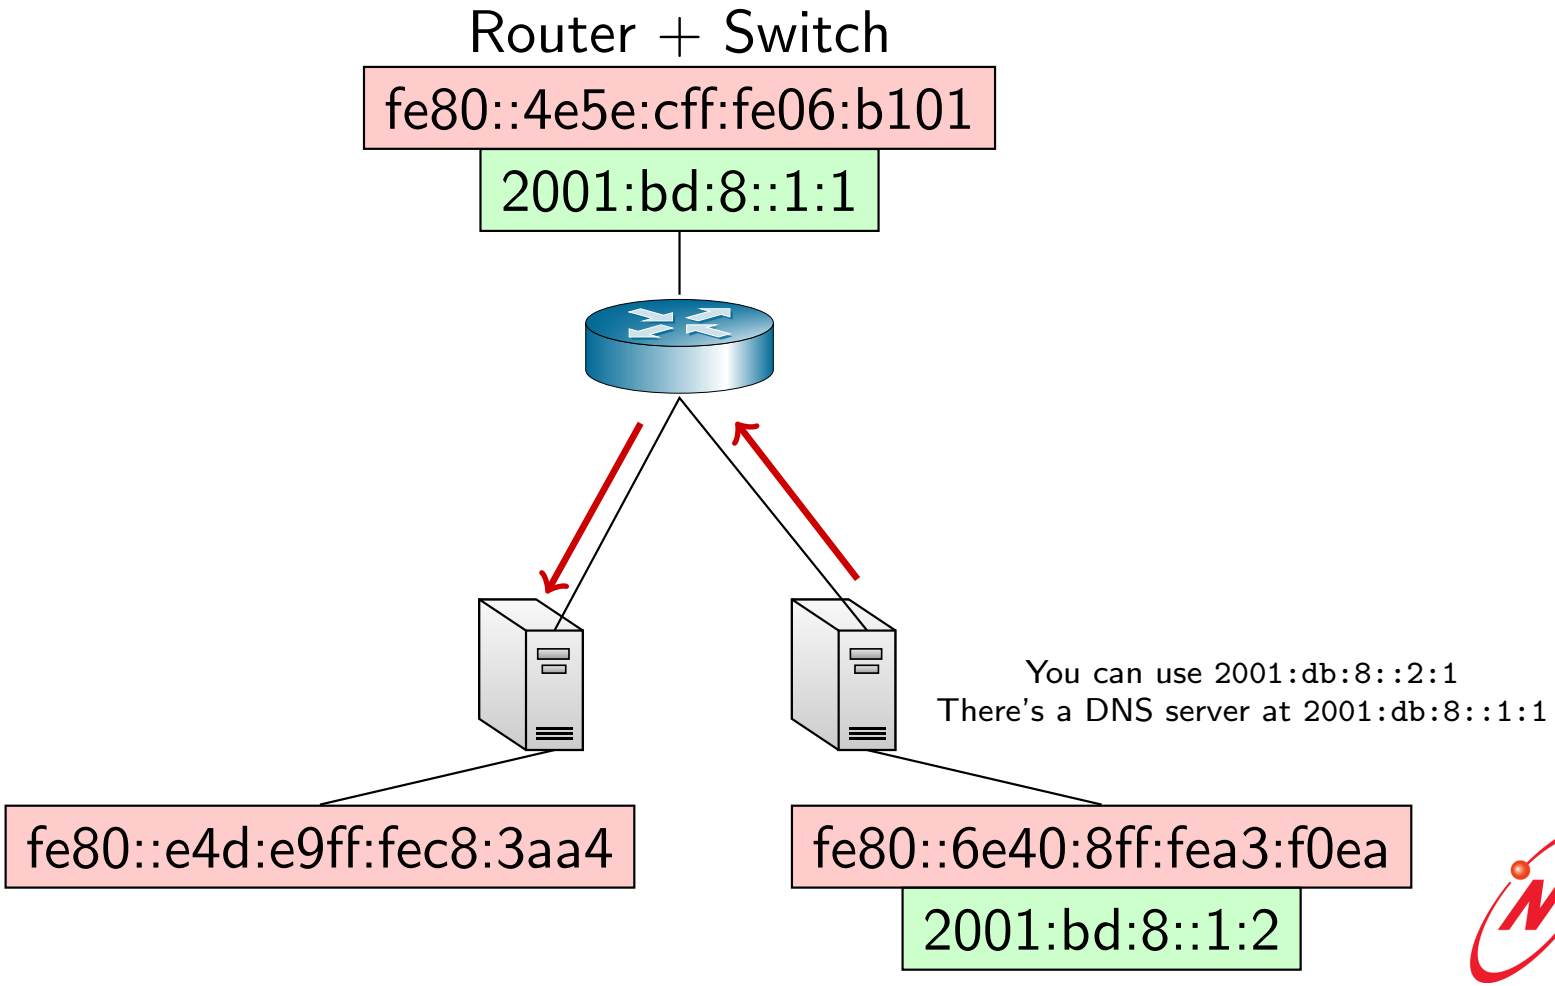

NetBurner
Networking in One Day!

Getting Global: DHCPv6

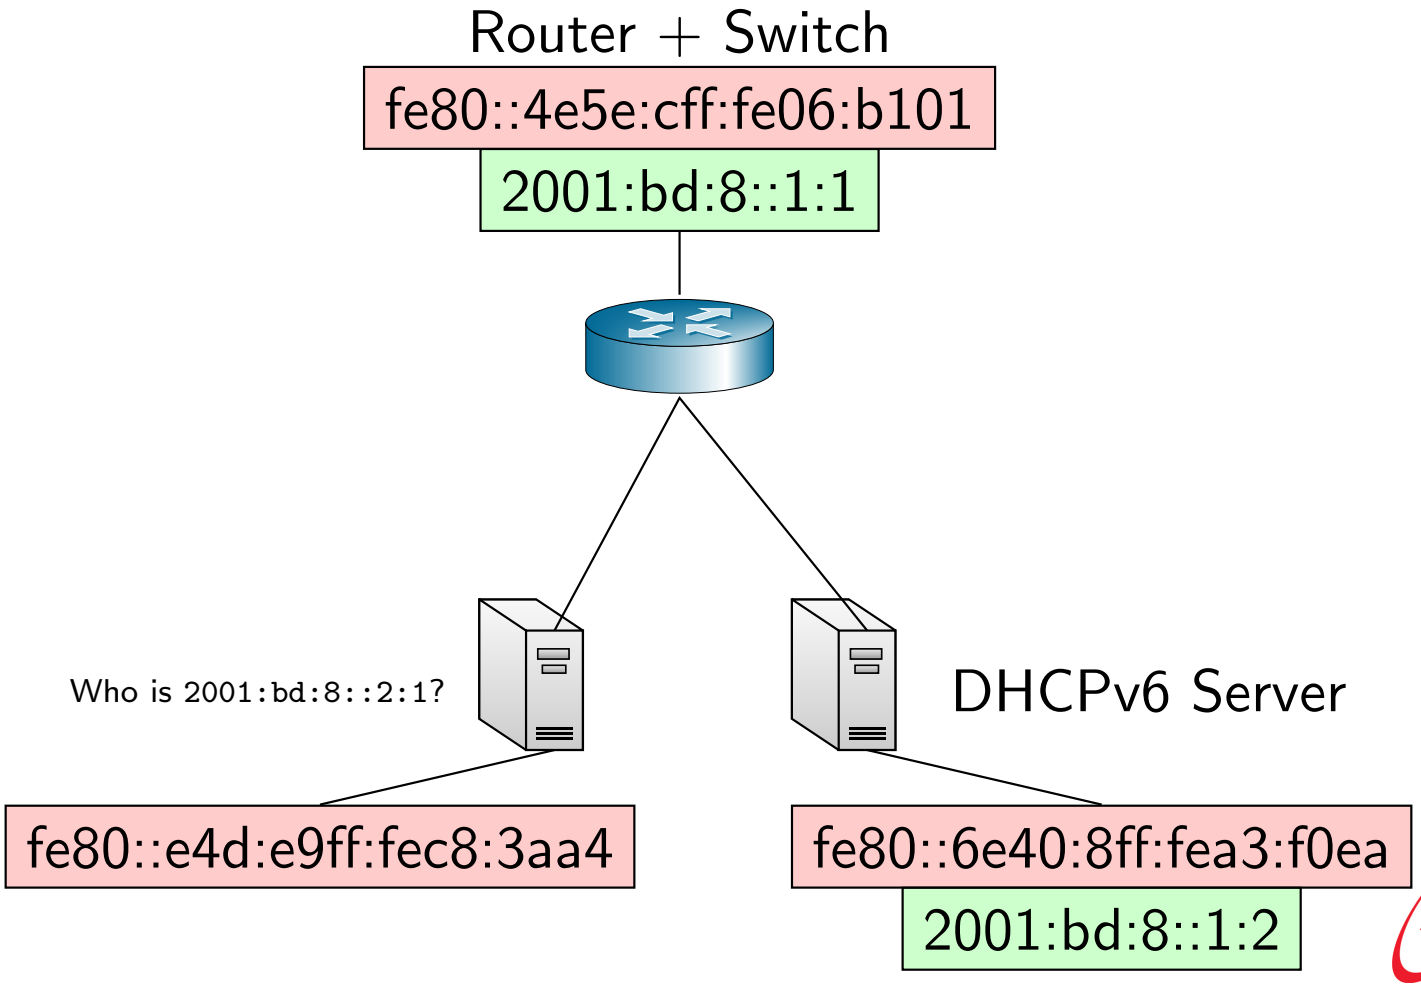

NetBurner
Networking in One Day!

Getting Global: DHCPv6

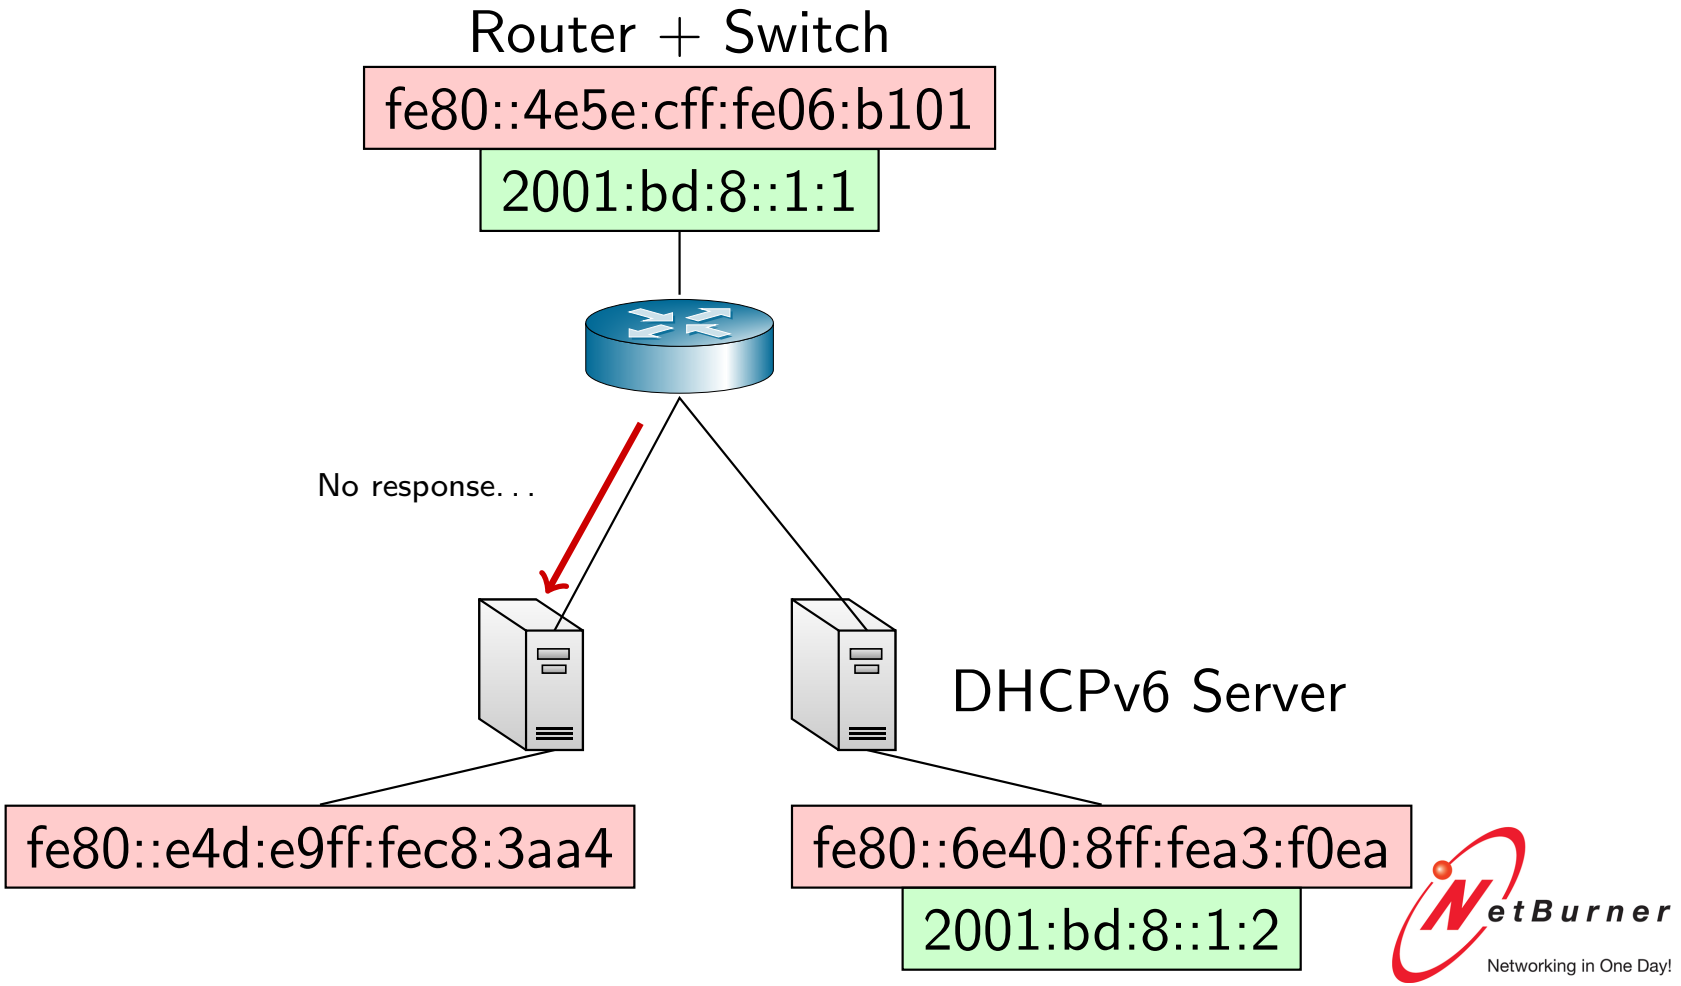

NetBurner
Networking in One Day!

Getting Global: DHCPv6

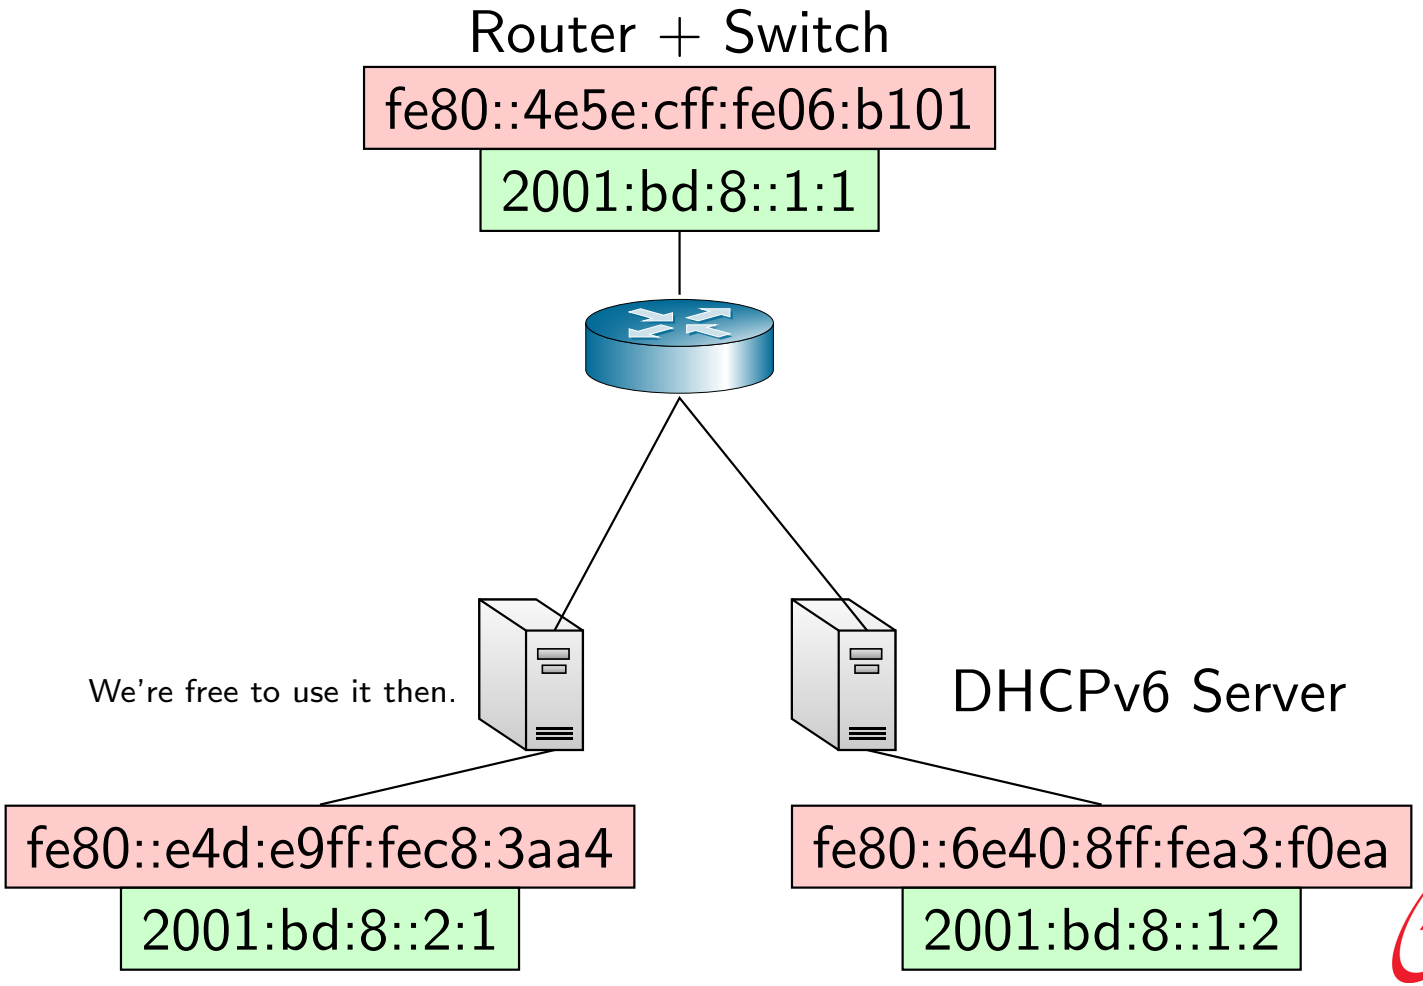

NetBurner
Networking in One Day!

We're going manual!

- Sometimes really want the device to have a specific address (i.e. 'DEAD:BEEF::BA5E:BA11').
- Theres no router or DHCPv6 server to hand out that address.
- So you need to assign it statically.
- We perform Duplicate Address Detection (DAD), and there we have it. A static address.

NetBurner
Networking in One Day!

Setting Static Address

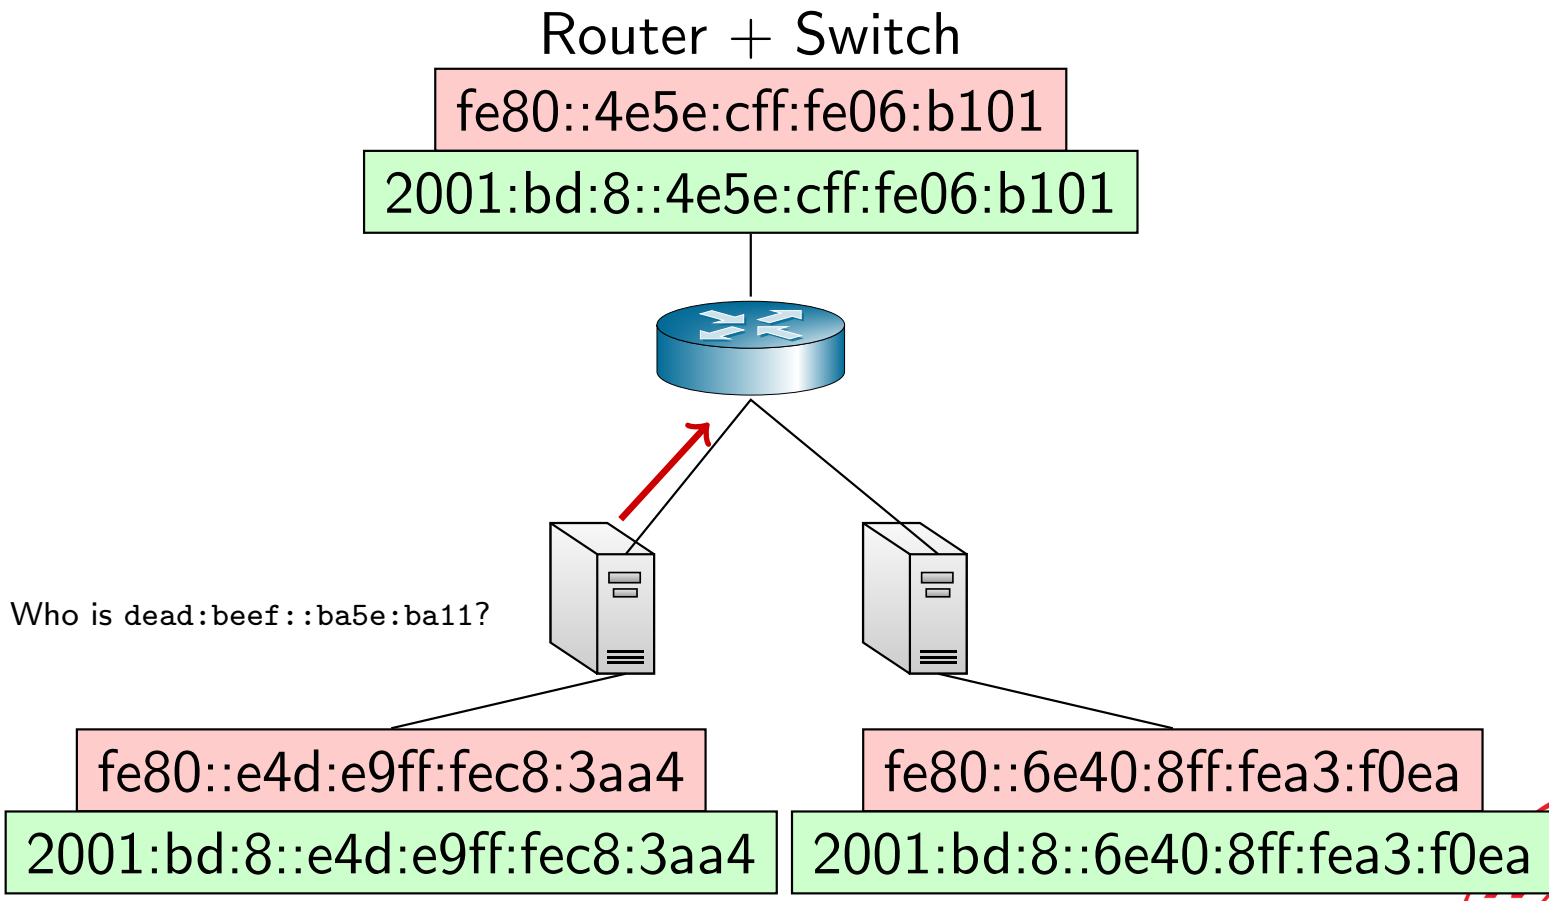

Setting Static Address

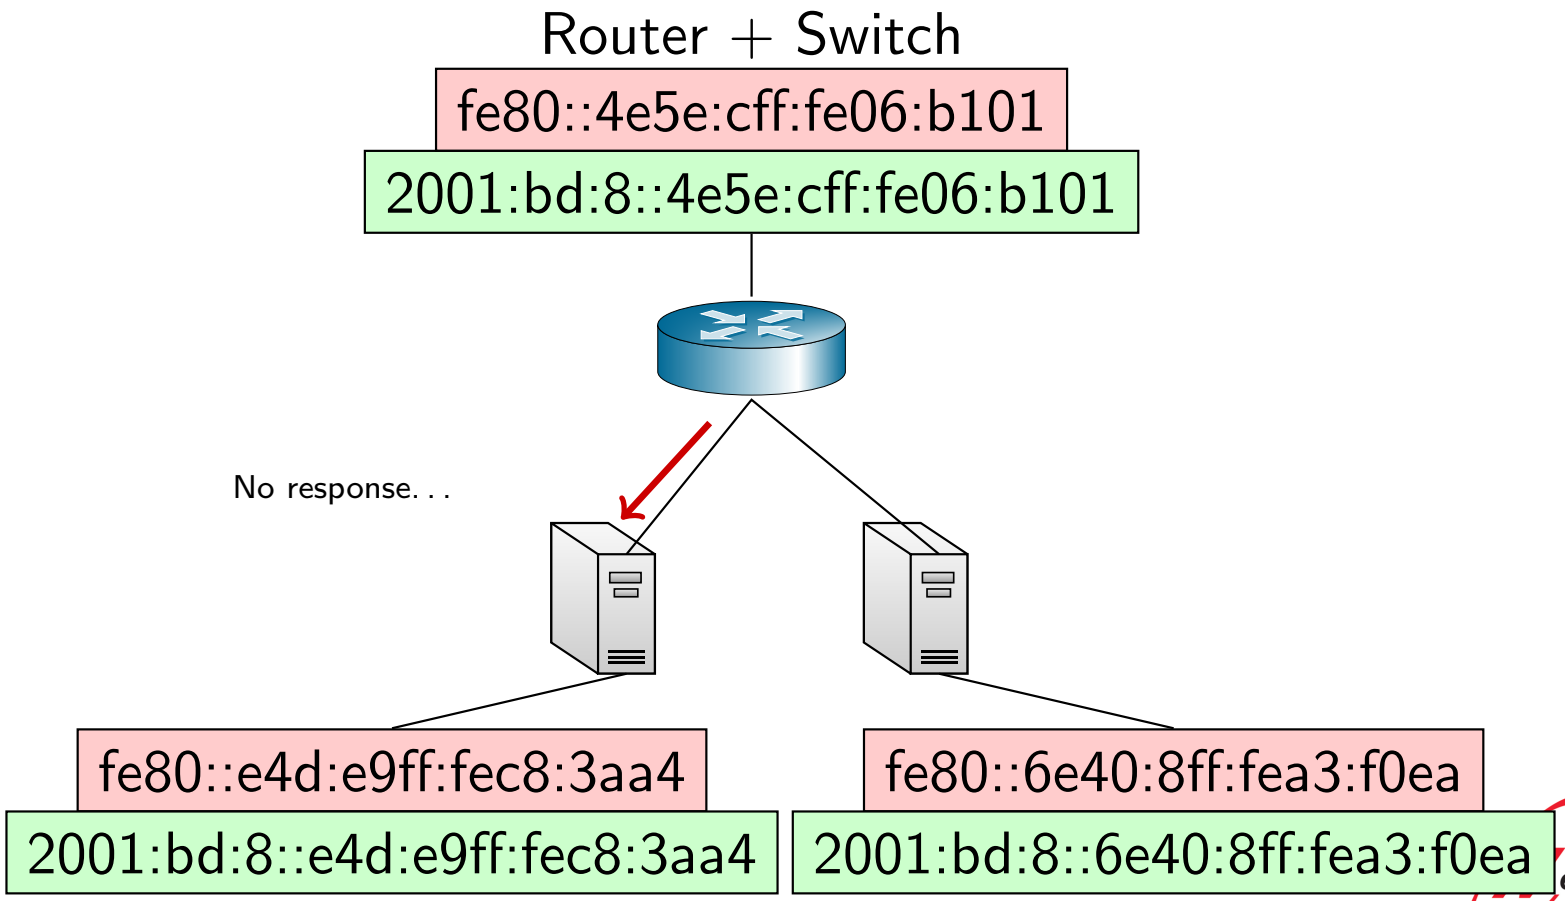

Setting Static Address

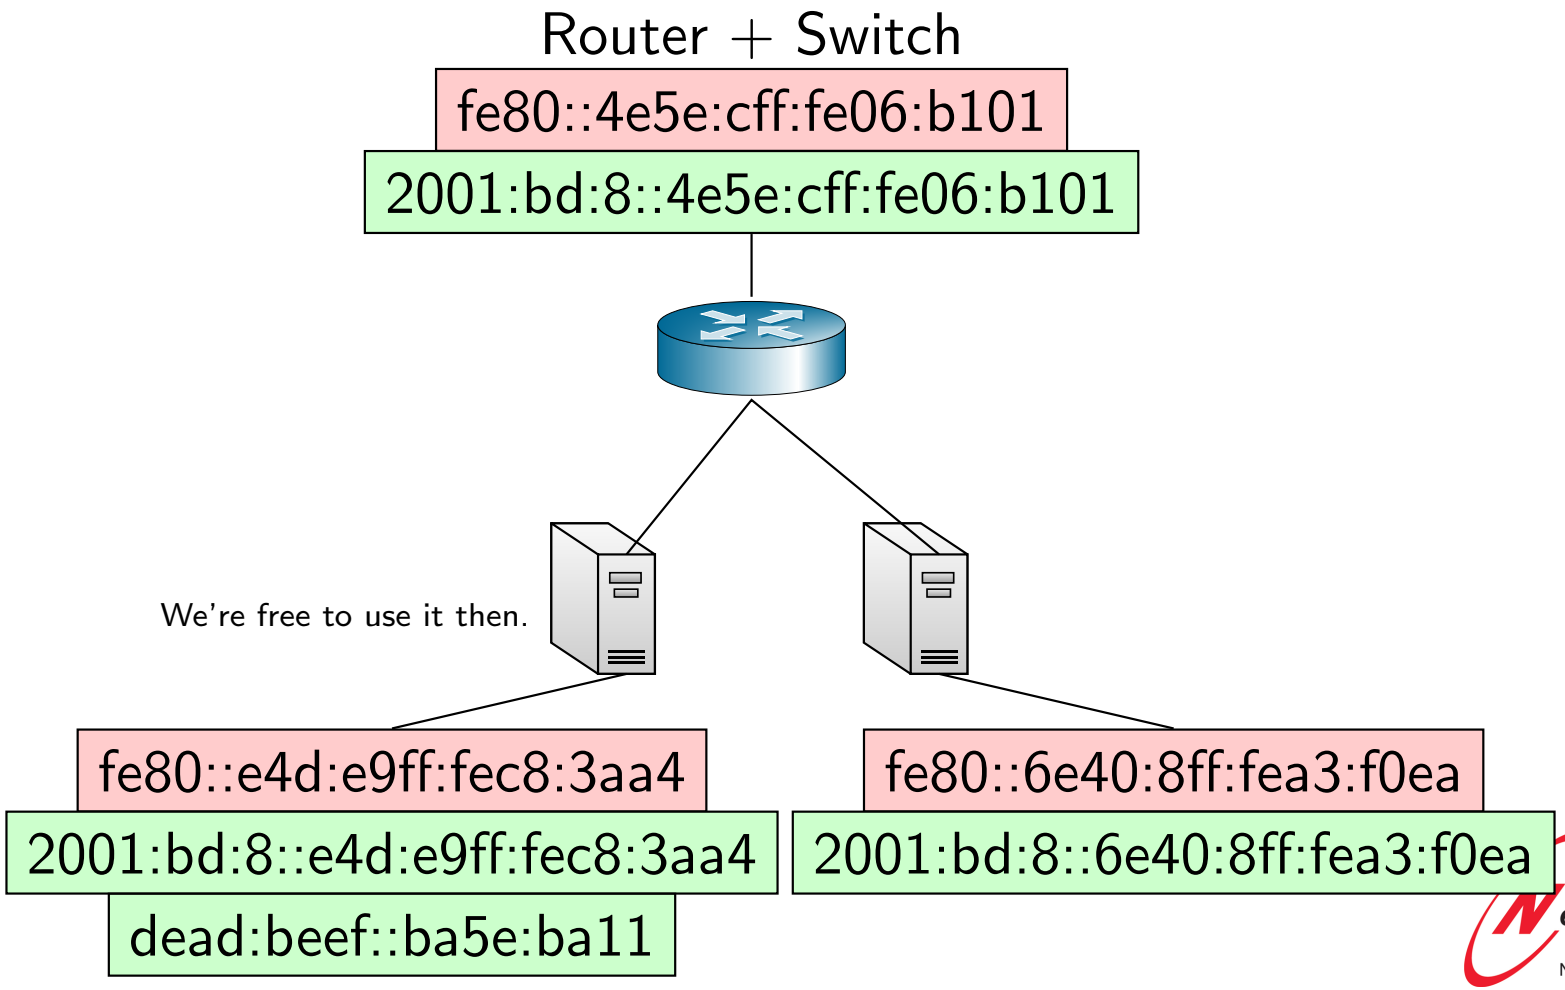

NetBurner
Networking in One Day!

How it goes wrong

Why can't I talk to my device?

IPv4: General

- The device is on a different subnet than you.
- There's an address conflict and some other device says it's at that address.
- The device has no address.

NetBurner
Networking in One Day!

IPv4: Different Subnet

Cause:

- At least one end of the communication chose the wrong address.

IPv4: Address Conflict

Cause:

- Someone chose poorly.

Solution:

- If mixed static/DHCP, move statics out of DHCP pool/adjust DHCP pool.

NetBurner
Networking in One Day!

IPv4: No Address

Cause by intended address scheme:

- DHCP: No DHCP server available or DHCP server is out of leases.
- AutoIP: We're still negotiating or Descartes' Demon is on this network...
- Static: No address was assigned.

Solutions:

- DHCP/Static: Assign an address manually.
- AutoIP: Wait. Hope. Pray.

NetBurner
Networking in One Day!

IPV6: General

- There's an address conflict.
- You're trying to connect to the web server on a link local address from OS X or BSD.
- Connecting to web server from not BSD or have already included scope ID.

IPv6: Address Conflict

Cause:

- There's a DHCPv6 Server doing managed address assignment and a router with the Autoconfig bit set for the same address space (i.e. prefix).

Solution:

- "Don't do that." -Rob Pike

NetBurner
Networking in One Day!

IPv6: Connecting from BSD Stack

Cause:

- The BSD IPv6 stack is dumber than a box of rocks.
- Tracks scope ID (aka interface to use) for link-local by stuffing into remote address.
- Therefore, scope ID required in browser link-local targets.

Example:

- Remote device is at `fe80::e4d:e9ff:fec8:3aa4`
- `en0` interface on Macbook has scope id `0xa`.
- Correct link for navigating is http://[fe80:**a**::e4d:e9ff:fec8:3aa4]/.

NetBurner
Networking in One Day!

IPv6: Browser issue not BSD scope ID

Cause:

- IPv6 addresses must be entered in brackets
  - E.G. http://fe80::6e40:8ff:fea3:f0ea/
    $\Rightarrow$ http://[fe80::6e40:8ff:fea3:f0ea]/
- This is for disambiguation between IPv6 addresses and specifying the port number after a domain.

NetBurner
Networking in One Day!